

Meta-Kernelization with Structural Parameters^{*}

Robert Ganian, Friedrich Slivovsky, and Stefan Szeider

Institute of Information Systems, Vienna University of Technology, Vienna, Austria
rganian@gmail.com, fslivovsky@gmail.com, stefan@szeider.net

Abstract. Meta-kernelization theorems are general results that provide polynomial kernels for large classes of parameterized problems. The known meta-kernelization theorems, in particular the results of Bodlaender et al. (FOCS'09) and of Fomin et al. (FOCS'10), apply to optimization problems parameterized by *solution size*. We present meta-kernelization theorems that use a *structural parameters* of the input and not the solution size. Let \mathcal{C} be a graph class. We define the \mathcal{C} -cover number of a graph to be the smallest number of modules the vertex set can be partitioned into such that each module induces a subgraph that belongs to the class \mathcal{C} .

We show that each graph problem that can be expressed in Monadic Second Order (MSO) logic has a polynomial kernel with a linear number of vertices when parameterized by the \mathcal{C} -cover number for any fixed class \mathcal{C} of bounded rank-width (or equivalently, of bounded clique-width, or bounded Boolean width). Many graph problems such as INDEPENDENT DOMINATING SET, c -COLORING, and c -DOMATIC NUMBER are covered by this meta-kernelization result.

Our second result applies to MSO expressible optimization problems, such as MINIMUM VERTEX COVER, MINIMUM DOMINATING SET, and MAXIMUM CLIQUE. We show that these problems admit a polynomial annotated kernel with a linear number of vertices.

1 Introduction

Kernelization is an algorithmic technique that has become the subject of a very active field in parameterized complexity, see, e.g., the references in [12,19,25]. Kernelization can be considered as a *preprocessing with performance guarantee* that reduces an instance of a parameterized problem in polynomial time to a decision-equivalent instance, the *kernel*, whose size is bounded by a function of the parameter alone [12,19,15]; if the reduced instance is an instance of a different problem, then it is called a *bikernel*. Once a kernel or bikernel is obtained, the time required to solve the original instance is bounded by a function of the parameter and therefore independent of the input size. Consequently one aims at (bi)kernels that are as small as possible.

Every fixed-parameter tractable problem admits a kernel, but the size of the kernel can have an exponential or even non-elementary dependence on the parameter [14]. Thus research on kernelization is typically concerned with the question of whether a fixed-parameter tractable problem under consideration admits a small, and in particular a *polynomial*, kernel. For instance, the parameterized MINIMUM VERTEX COVER problem (does a given graph have a vertex cover consisting of k vertices?) admits a polynomial kernel containing at most $2k$ vertices. There are many fixed-parameter tractable problems for which no polynomial kernels are known. Recently, theoretical tools have been developed to provide strong theoretical evidence that certain fixed-parameter tractable problems do not admit polynomial kernels [3]. In particular, these techniques can be applied to a wide range of graph problems parameterized by treewidth and other width parameters such as clique-width, or rank-width. Thus, in order to get polynomial kernels, structural parameters have been suggested that are somewhat weaker than treewidth, including the vertex cover number, max-leaf number, and neighborhood diversity [13,21]. The general aim is to find a parameter that admits a polynomial kernel while being as general as possible.

We extend this line of research by using results from modular decompositions and rank-width to introduce new structural parameters for which large classes of problems have polynomial kernels. Specifically, we study the *rank-width- d cover number*, which is a special case of a \mathcal{C} -cover number (see Section 3 for definitions). We establish the following result which is an important prerequisite for our kernelization results.

Theorem 1. *For every constant d , a smallest rank-width- d cover of a graph can be computed in polynomial time.*

^{*} Research supported by the European Research Council (ERC), project COMPLEX REASON 239962.

Hence, for graph problems parameterized by rank-width- d cover number, we can always compute the parameter in polynomial time. The proof of Theorem 1 relies on a combinatorial property of modules of bounded rank-width that amounts to a variant of partitivity [8].

Our kernelization results take the shape of *algorithmic meta-theorems*, stated in terms of the evaluation of formulas of monadic second order logic (MSO) on graphs. Monadic second order logic over graphs extends first order logic by variables that may range over sets of vertices (sometimes referred to as MSO_1 logic). Specifically, for an MSO formula φ , our first meta-theorem applies to all problems of the following shape, which we simply call *MSO model checking* problems.

MSO-MC $_{\varphi}$

Instance: A graph G .

Question: Does $G \models \varphi$ hold?

Many NP-hard graph problems can be naturally expressed as MSO model checking problems, for instance INDEPENDENT DOMINATING SET, c -COLORING, and c -DOMATIC NUMBER.

Theorem 2. *Let \mathcal{C} be a graph class of bounded rank-width. Every MSO model checking problem, parameterized by the \mathcal{C} -cover number of the input graph, has a polynomial kernel with a linear number of vertices.*

While MSO model checking problems already capture many important graph problems, there are some well-known optimization problems on graphs that cannot be captured in this way, such as MINIMUM VERTEX COVER, MINIMUM DOMINATING SET, and MAXIMUM CLIQUE. Many such optimization graph problems can be stated in the following way. Let $\varphi = \varphi(X)$ be an MSO formula with one free set variable X and $\diamond \in \{\leq, \geq\}$.

MSO-OPT $_{\varphi}^{\diamond}$

Instance: A graph G and an integer $r \in \mathbb{N}$.

Question: Is there a set $S \subseteq V(G)$ such that $G \models \varphi(S)$ and $|S| \diamond r$?

We call problems of this form *MSO optimization problems*. MSO optimization problems form a large fragment of the so-called *LinMSO* problems [2]. There are dozens of well-known graph problems that can be expressed as MSO optimization problems.

We establish the following result.

Theorem 3. *Let \mathcal{C} be a graph class of bounded rank-width. Every MSO optimization problem, parameterized by the \mathcal{C} -cover number of the input graph, has a polynomial bikernel with a linear number of vertices.*

In fact, the obtained bikernel is an instance of an annotated variant of the original MSO optimization problem [1]. Hence, Theorem 3 provides a polynomial kernel for an annotated version of the original MSO optimization problem.

For obtaining the kernel for MSO model checking problems we proceed as follows. First we compute a smallest rank-width- d cover of the input graph G in polynomial time. Second, we compute for each module a small representative of constant size. Third, we replace each module with a constant size module, which results in the kernel. For the MSO optimization problems we proceed similarly. However, in order to represent a possibly large module with a small module of constant size, we need to keep the information how much a solution projected on a module contributes to the full solution. We provide this information by means of annotations to the kernel.

We would like to point out that a class of graphs has bounded rank-width iff it has bounded clique-width iff it has bounded Boolean-width [6]. Hence, we could have equivalently stated the theorems in terms of clique-width or Boolean width. Furthermore we would like to point out that the theorems hold also for some classes \mathcal{C} where we do not know whether \mathcal{C} can be recognized in polynomial time, and where we do not know how to compute the partition in polynomial time. For instance, the theorems hold if \mathcal{C} is a graph class of bounded clique-width (it is not known whether graphs of clique-width at most 4 can be recognized in polynomial time).

2 Preliminaries

The set of natural numbers (that is, positive integers) will be denoted by \mathbb{N} . For $i \in \mathbb{N}$ we write $[i]$ to denote the set $\{1, \dots, i\}$.

Graphs. We will use standard graph theoretic terminology and notation (cf. [10]). A *module* of a graph $G = (V, E)$ is a nonempty set $X \subseteq V$ such that for each vertex $v \in V \setminus X$ it holds that either no element of X is a neighbor of v or every element of X is a neighbor of v . We say two modules $X, Y \subseteq V$ are *adjacent* if there are vertices $x \in X$ and $y \in Y$ such that x and y are adjacent. A *modular partition* of a graph G is a partition $\{U_1, \dots, U_k\}$ of its vertex set such that U_i is a module of G for each $i \in [k]$.

Monadic Second-Order Logic on Graphs. We assume that we have an infinite supply of individual variables, denoted by lowercase letters x, y, z , and an infinite supply of set variables, denoted by uppercase letters X, Y, Z . Formulas of *monadic second-order logic* (MSO) are constructed from atomic formulas $E(x, y)$, $X(x)$, and $x = y$ using the connectives \neg (negation), \wedge (conjunction) and existential quantification $\exists x$ over individual variables as well as existential quantification $\exists X$ over set variables. Individual variables range over vertices, and set variables range over sets of vertices. The atomic formula $E(x, y)$ expresses adjacency, $x = y$ expresses equality, and $X(x)$ expresses that vertex x is in the set X . From this, we define the semantics of monadic second-order logic in the standard way (this logic is sometimes called MSO_1).

Free and bound variables of a formula are defined in the usual way. A *sentence* is a formula without free variables. We write $\varphi(X_1, \dots, X_n)$ to indicate that the set of free variables of formula φ is $\{X_1, \dots, X_n\}$. If $G = (V, E)$ is a graph and $S_1, \dots, S_n \subseteq V$ we write $G \models \varphi(S_1, \dots, S_n)$ to denote that φ holds in G if the variables x_i are interpreted by the vertices v_i and the variables X_j are interpreted by the sets S_j ($i \in [n], j \in [m]$).

We review MSO *types* and *games* roughly following the presentation in [22]. The *quantifier rank* of an MSO formula φ is defined as the nesting depth of quantifiers in φ . For non-negative integers q and l , let $\text{MSO}_{q,l}$ consist of all MSO formulas of quantifier rank at most q with free set variables in $\{X_1, \dots, X_l\}$.

Let $\varphi = \varphi(X_1, \dots, X_l)$ and $\psi = \psi(X_1, \dots, X_l)$ be MSO formulas. We say φ and ψ are *equivalent*, written $\varphi \equiv \psi$, if for all graphs G and $U_1, \dots, U_l \subseteq V(G)$, $G \models \varphi(U_1, \dots, U_l)$ if and only if $G \models \psi(U_1, \dots, U_l)$. Given a set F of formulas, let F/\equiv denote the set of equivalence classes of F with respect to \equiv . The following statement has a straightforward proof using normal forms (see Theorem 7.5 in [22] for details).

Fact 1. *Let q and l be non-negative integers. The set $\text{MSO}_{q,l}/\equiv$ is finite, and given q and l one can effectively compute a system of representatives of $\text{MSO}_{q,l}/\equiv$.*

We will assume that for any pair of non-negative integers q and l the system of representatives of $\text{MSO}_{q,l}/\equiv$ given by Fact 1 is fixed.

Definition 4 (MSO Type). *Let q, l be non-negative integers. For a graph G and an l -tuple \mathbf{U} of sets of vertices of G , we define $\text{type}_q(G, \mathbf{U})$ as the set of formulas $\varphi \in \text{MSO}_{q,l}$ such that $G \models \varphi(\mathbf{U})$. We call $\text{type}_q(G, \mathbf{U})$ the MSO rank- q type of \mathbf{U} in G .*

It follows from Fact 1 that up to logical equivalence, every type contains only finitely many formulas. This allows us to represent types using MSO formulas as follows.

Lemma 5. *Let q and l be non-negative integer constants, let G be a graph, and let \mathbf{U} be an l -tuple of sets of vertices of G . One can effectively compute a formula $\Phi \in \text{MSO}_{q,l}$ such that for any graph G' and any l -tuple \mathbf{U}' of sets of vertices of G' we have $G' \models \Phi(\mathbf{U}')$ if and only if $\text{type}_q(G, \mathbf{U}) = \text{type}_q(G', \mathbf{U}')$. Moreover, if $G \models \varphi(\mathbf{U})$ can be decided in polynomial time for any fixed $\varphi \in \text{MSO}_{q,l}$ then Φ can be computed in time polynomial in $|V(G)|$.*

Proof. Let R be a system of representatives of $\text{MSO}_{q,l}/\equiv$ given by Fact 1. Because q and l are constant, we can consider both the cardinality of R and the time required to compute it as constants. Let $\Phi \in \text{MSO}_{q,l}$ be the formula defined as $\Phi = \bigwedge_{\varphi \in S} \varphi \wedge \bigwedge_{\varphi \in R \setminus S} \neg \varphi$, where $S = \{\varphi \in R : G \models \varphi(\mathbf{U})\}$. We can compute Φ by deciding $G \models \varphi(\mathbf{U})$ for each $\varphi \in R$. Since the number of formulas in R is a constant, this can be done in polynomial time if $G \models \varphi(\mathbf{U})$ can be decided in polynomial time for any fixed $\varphi \in \text{MSO}_{q,l}$.

Let G' be an arbitrary graph and \mathbf{U}' an l -tuple of subsets of $V(G')$. We claim that $\text{type}_q(G, \mathbf{U}) = \text{type}_q(G', \mathbf{U}')$ if and only if $G' \models \Phi(\mathbf{U}')$. Since $\Phi \in \text{MSO}_{q,l}$ the forward direction is trivial. For the converse, assume $\text{type}_q(G, \mathbf{U}) \neq \text{type}_q(G', \mathbf{U}')$. First suppose $\varphi \in \text{type}_q(G, \mathbf{U}) \setminus \text{type}_q(G', \mathbf{U}')$. The set R is a system of representatives of $\text{MSO}_{q,l}/\equiv$, so there has to be a $\psi \in R$ such that $\psi \equiv \varphi$. But $G' \models \Phi(\mathbf{U}')$ implies $G' \models \psi(\mathbf{U}')$ by construction of Φ and thus $G' \models \varphi(\mathbf{U}')$, a contradiction. Now suppose $\varphi \in \text{type}_q(G', \mathbf{U}') \setminus \text{type}_q(G, \mathbf{U})$. An analogous argument proves that there has to be a $\psi \in R$ such that $\psi \equiv \varphi$ and $G' \models \neg \psi(\mathbf{U}')$. It follows that $G' \not\models \varphi(\mathbf{U}')$, which again yields a contradiction. \square

Definition 6 (Partial isomorphism). Let G, G' be graphs, and let $\mathbf{V} = (V_1, \dots, V_l)$ and $\mathbf{U} = (U_1, \dots, U_l)$ be tuples of sets of vertices with $V_i \subseteq V(G)$ and $U_i \subseteq V(G')$ for each $i \in [l]$. Let $\mathbf{v} = (v_1, \dots, v_m)$ and $\mathbf{u} = (u_1, \dots, u_m)$ be tuples of vertices with $v_i \in V(G)$ and $u_i \in V(G')$ for each $i \in [m]$. Then (\mathbf{v}, \mathbf{u}) defines a partial isomorphism between (G, \mathbf{V}) and (G', \mathbf{U}) if the following conditions hold:

– For every $i, j \in [m]$,

$$v_i = v_j \Leftrightarrow u_i = u_j \text{ and } v_i v_j \in E(G) \Leftrightarrow u_i u_j \in E(G').$$

– For every $i \in [m]$ and $j \in [l]$,

$$v_i \in V_j \Leftrightarrow u_i \in U_j.$$

Definition 7. Let G and G' be graphs, and let \mathbf{V}_0 be a k -tuple of subsets of $V(G)$ and let \mathbf{U}_0 be a k -tuple of subsets of $V(G')$. Let q be a non-negative integer. The q -round MSO game on G and G' starting from $(\mathbf{V}_0, \mathbf{U}_0)$ is played as follows. The game proceeds in rounds, and each round consists of one of the following kinds of moves.

- **Point move** The spoiler picks a vertex in either G or G' ; the duplicator responds by picking a vertex in the other graph.
- **Set move** The spoiler picks a subset of $V(G)$ or a subset of $V(G')$; the duplicator responds with a subset of the vertex set of the other graph.

Let $v_1, \dots, v_m \in V(G)$ and $u_1, \dots, u_m \in V(G')$ be the point moves played in the q -round game, and let $V_1, \dots, V_l \subseteq V(G)$ and $U_1, \dots, U_l \subseteq V(G')$ be the set moves played in the q -round game, so that $l + m = q$ and moves belonging to same round have the same index. Then the duplicator wins the game if (\mathbf{v}, \mathbf{u}) is a partial isomorphism of $(G, \mathbf{V}_0 \mathbf{V})$ and $(G', \mathbf{U}_0 \mathbf{U})$. If duplicator has a winning strategy, we write $(G, \mathbf{V}_0) \equiv_q^{MSO} (G', \mathbf{U}_0)$.

Theorem 8 ([22], Theorem 7.7). Given two graphs G and G' and two l -tuples $\mathbf{V}_0, \mathbf{U}_0$ of sets of vertices of G and G' , we have

$$\text{type}_q(G, \mathbf{V}_0) = \text{type}_q(G, \mathbf{U}_0) \Leftrightarrow (G, \mathbf{V}_0) \equiv_q^{MSO} (G', \mathbf{U}_0).$$

Fixed-Parameter Tractability and Kernels. A parameterized problem P is a subset of $\Sigma^* \times \mathbb{N}$ for some finite alphabet Σ . For a problem instance $(x, k) \in \Sigma^* \times \mathbb{N}$ we call x the main part and k the parameter. A parameterized problem P is *fixed-parameter tractable* (FPT) if a given instance (x, k) can be solved in time $O(f(k) \cdot p(|x|))$ where f is an arbitrary computable function of k and p is a polynomial in the input size $|x|$.

A *bikernelization* for a parameterized problem $P \subseteq \Sigma^* \times \mathbb{N}$ into a parameterized problem $Q \subseteq \Sigma^* \times \mathbb{N}$ is an algorithm that, given $(x, k) \in \Sigma^* \times \mathbb{N}$, outputs in time polynomial in $|x| + k$ a pair $(x', k') \in \Sigma^* \times \mathbb{N}$ such that (i) $(x, k) \in P$ if and only if $(x', k') \in Q$ and (ii) $|x'| + k' \leq g(k)$, where g is an arbitrary computable function. The reduced instance (x', k') is the *bikernel*. If $P = Q$, the reduction is called a *kernelization* and (x', k') a *kernel*. The function g is called the *size* of the (bi)kernel, and if g is a polynomial then we say that P admits a *polynomial (bi)kernel*.

It is well known that every fixed-parameter tractable problem admits a generic kernel, but the size of this kernel can have an exponential or even non-elementary dependence on the parameter [11]. Since recently there have been workable tools available for providing strong theoretical evidence that certain parameterized problems do not admit a polynomial kernel [3,23].

Rank-width The graph invariant rank-width was introduced by Oum and Seymour [24] with the original intent of investigating the graph invariant clique-width. It later turned out that rank-width itself is a useful parameter, with several advantages over clique-width.

A set function $f : 2^M \rightarrow \mathbb{Z}$ is called *symmetric* if $f(X) = f(M \setminus X)$ for all $X \subseteq M$. For a symmetric function $f : 2^M \rightarrow \mathbb{Z}$ on a finite set M , a *branch-decomposition* of f is a pair (T, μ) where T tree of maximum degree 3 and $\mu : M \rightarrow \{t : t \text{ is a leaf of } T\}$ is a bijective function. For an edge e of T , the connected components of $T \setminus e$ induce a bipartition (X, Y) of the set of leaves of T . The *width* of an edge e of a branch-decomposition (T, μ) is $f(\mu^{-1}(X))$. The *width* of (T, μ) is the maximum width over all edges of T . The *branch-width* of f




Fig. 1. Relationship between graph invariants: the vertex cover number (vcn), the neighborhood diversity (nd), the rank-width- d cover number ($rw_{c,d}$), the rank-width (rw), and the treewidth (tw). An arrow from A to B indicates that for any graph class for which B is bounded also A is bounded.

is the minimum width over all branch-decompositions of f . If $|M| \leq 1$, then we define the branch-width of f as $f(\emptyset)$. A natural application of this definition is the branch-width of a graph, as introduced by Robertson and Seymour [?], where $M = E(G)$, and f the connectivity function of G .

There is, however, another interesting application of the aforementioned general notions, in which we consider the vertex set $V(G) = M$ of a graph G as the ground set.

For a graph G and $U, W \subseteq V(G)$, let $\mathbf{A}_G[U, W]$ denote the $U \times W$ -submatrix of the adjacency matrix over the two-element field $\text{GF}(2)$, i.e., the entry $a_{u,w}$, $u \in U$ and $w \in W$, of $\mathbf{A}_G[U, W]$ is 1 if and only if $\{u, w\}$ is an edge of G . The *cut-rank* function ρ_G of a graph G is defined as follows: For a bipartition (U, W) of the vertex set $V(G)$, $\rho_G(U) = \rho_G(W)$ equals the rank of $\mathbf{A}_G[U, W]$ over $\text{GF}(2)$. A *rank-decomposition* and *rank-width* of a graph G is the branch-decomposition and branch-width of the cut-rank function ρ_G of G on $M = V(G)$, respectively.

Theorem 9 ([20]). *Let $k \in \mathbb{N}$ be a constant and $n \geq 2$. For an n -vertex graph G , we can output a rank-decomposition of width at most k or confirm that the rank-width of G is larger than k in time $O(n^3)$.*

Theorem 10 ([18]). *Let $d \in \mathbb{N}$ be a constant and let φ and $\psi = \psi(X)$ be fixed MSO formulas. Given a graph G with $rw(G) \leq d$, one can decide whether $G \models \varphi$ in polynomial time. Moreover, a set $S \subseteq V(G)$ of minimum (maximum) cardinality such that $G \models \psi(S)$ can be found in polynomial time, if one exists.*

3 Rank-width Covers

Let G^1 be the trivial single-vertex graph, and let \mathcal{C} be a graph class such that $G^1 \in \mathcal{C}$. We define a \mathcal{C} -cover of G as a modular partition $\{U_1, \dots, U_k\}$ of $V(G)$ such that the induced subgraph $G[U_i]$ belongs to the class \mathcal{C} for each $i \in [k]$. Accordingly, the \mathcal{C} -cover number of G is the size of a smallest \mathcal{C} -cover of G .

Of special interest to us are the classes \mathcal{R}_d of graphs of rank-width at most d . We call the \mathcal{R}_d -cover number also the *rank-width- d cover number*. If \mathcal{C} is the class of complete and edgeless graphs, then the \mathcal{C} -cover number equals the neighborhood diversity [21], and clearly $\mathcal{C} \subsetneq \mathcal{R}_1$. Figure 1 shows the relationship between the rank-width- d cover number and some other graph invariants.

We state some further properties of rank-width- d covers.

Proposition 11. *Let vcn , nd , and rw denote the vertex cover number, the neighborhood diversity, and the rank-width of a graph G , respectively. Then the following (in)equalities hold for any $d \in \mathbb{N}$:*

1. $rw_{c,d}(G) \leq nd(G) \leq 2^{vcn(G)}$,
2. if $d \geq rw(G)$, then $|rw_{c,d}(G)| = 1$.

Proof. (1) The neighborhood diversity of a graph is also a rank-width-1 cover. The neighborhood diversity is known to be upper-bounded by $2^{vcn(G)}$ [21]. (2) follows immediately from the definition of rank-width- d covers. \square

3.1 Finding the Cover

Next we state several properties of modules of graphs. These will be used to obtain a polynomial algorithm for finding smallest rank-width- d covers.

The *symmetric difference* of sets A, B is $A \triangle B = (A \setminus B) \cup (B \setminus A)$. Sets A, B *overlap* if $A \cap B \neq \emptyset$ but neither $A \subseteq B$ nor $B \subseteq A$.

Definition 12. Let $\mathcal{S} \subseteq 2^S$ be a family of subsets of a set S . We call \mathcal{S} *partitive* if it satisfies the following properties:

1. $S \in \mathcal{S}$, $\emptyset \notin \mathcal{S}$, and $\{x\} \in \mathcal{S}$ for each $x \in S$.
2. For every pair of overlapping subsets $A, B \in \mathcal{S}$, the sets $A \cup B$, $A \cap B$, $A \triangle B$, $A \setminus B$, and $B \setminus A$ are contained in \mathcal{S} .

Theorem 13 ([8]). The family of modules of a graph G is partitive.

Lemma 14 ([5]). Let G be a graph and $x, y \in V(G)$. There is a unique minimal (with respect to set inclusion) module M of G such that $x, y \in M$, and M can be computed in time $O(|V(G)|^2)$.

Lemma 15. Let $d \in \mathbb{N}$ be a constant. Let G be a graph and let M_1, M_2 be modules of G such that $M_1 \cap M_2 \neq \emptyset$ and $\max(rw(G[M_1]), rw(G[M_2])) \leq d$. Then $M_1 \cup M_2$ is a module of G and $rw(G[M_1 \cup M_2]) \leq d$.

Proof. If $M_1 \subseteq M_2$ or $M_2 \subseteq M_1$ the result is immediate. Suppose M_1 and M_2 overlap and let $M_{11} = M_1 \setminus M_2$, $M_{22} = M_2 \setminus M_1$, and $M_{12} = M_1 \cap M_2$. It follows from Theorem 13 that these sets are modules of G . Let $v_{11} \in M_{11}$, $v_{22} \in M_{22}$, and $v_{12} \in M_{12}$. We show that $rw(G[M_1 \cup M_2]) \leq d$. By assumption, both $G[M_1]$ and $G[M_2]$ have rank-width at most d . Since rank-width is preserved by taking induced subgraphs, the graphs $G_{11} = G[M_{11} \cup \{v_{12}\}]$, $G_{12} = G[M_{12} \cup \{v_{22}\}]$, and $G_{22} = G[M_{22} \cup \{v_{12}\}]$ also have rank-width at most d . Let $\mathcal{T}_{11} = (T_{11}, \mu_{11})$, $\mathcal{T}_{12} = (T_{12}, \mu_{12})$, and $\mathcal{T}_{22} = (T_{22}, \mu_{22})$ be witnessing rank decompositions of G_{11} , G_{12} , and G_{22} , respectively.

We construct a rank decomposition $\mathcal{T} = (T, \mu)$ of $G[M_1 \cup M_2]$ as follows. Let l_{22} be the leaf (note that μ_{12} is bijective) of T_{12} such that $\mu_{12}(v_{22}) = l_{22}$. Moreover, let l_{12} and l'_{12} be the leaves of T_{11} and T_{22} such that $\mu_{11}(v_{12}) = l_{12}$ and $\mu_{22}(v_{12}) = l'_{12}$, respectively. We obtain T from T_{12} by adding disjoint copies of T_{11} and T_{22} and then identifying l_{22} with the copies of l_{12} and l'_{12} . Since T_{11} , T_{12} , and T_{22} are subcubic, so is T .

We define the mapping $\mu : M_1 \cup M_2 \rightarrow \{t : t \text{ is a leaf of } T\}$ by

$$\mu(v) = \begin{cases} \mu_{12}(v) & \text{if } v \in M_{12}, \\ c(\mu_{11}(v)) & \text{if } v \in M_{11}, \\ c(\mu_{22}(v)) & \text{otherwise,} \end{cases}$$

where c maps nodes in $T_{11} \cup T_{22}$ to their copies in T . The mappings μ_{11} , μ_{12} , and μ_{22} are bijections and c is injective, so μ is injective. By construction, the image of $M_1 \cup M_2$ under μ is the set of leaves of T , so μ is a bijection. Thus $\mathcal{T} = (T, \mu)$ is a rank decomposition of $G[M_1 \cup M_2]$.

We prove that the width of \mathcal{T} is at most d . Given a rank decomposition $\mathcal{T}^* = (T^*, \mu^*)$ and an edge $e \in T^*$, the connected components of $T^* \setminus \{e\}$ induce a bipartition (X, Y) of the leaves of T^* . We set $f : (T^*, e) \mapsto (\mu^{*-1}(X), \mu^{*-1}(Y))$. Take any edge e of T . There is a natural bijection β from the edges in T to the edges of $T_{11} \cup T_{12} \cup T_{22}$. Accordingly, we distinguish three cases for $e' = \beta(e)$:

1. $e' \in T_{11}$. Let $(U, W) = f(\mathcal{T}_{11}, e')$. Without loss of generality assume that $v_{12} \in W$. Then by construction of \mathcal{T} , we have $f(\mathcal{T}, e) = (U, W \cup M_2)$. Pick any $u \in U \subseteq M_{11}$ and $v \in M_2 \setminus W$. Since M_2 is a module of G with $v, v_{12} \in M_2$ but $u \notin M_2$ we have $\mathbf{A}_G(u, v) = \mathbf{A}_G(u, v_{12})$. As a consequence, $\mathbf{A}_G[U, W \cup M_2]$ can be obtained from $\mathbf{A}_G[U, W]$ by copying the column corresponding to v_{12} . This does not increase the rank of the matrix.
2. $e' \in T_{22}$. This case is symmetric to case 1, with M_{22} and M_1 taking the roles of M_{11} and M_2 , respectively.
3. $e' \in T_{12}$. Let $(U, W) = f(\mathcal{T}_{12}, e')$. Without loss of generality assume that $v_{22} \in W$. Then $f(\mathcal{T}, e) = (U, W \cup M_{11} \cup M_{22})$. Let $u \in U \subseteq M_{12}$ and $v \in M_{22}$. Since M_1 is a module and $u \in M_1$ but $v, v_{22} \notin M_1$, we must have $\mathbf{A}_G(u, v) = \mathbf{A}_G(u, v_{22})$, so one can simply copy the column corresponding to v_{22} . Now consider $w \in M_{11}$. Suppose $wu \in E(G)$. Since $u, v_{22} \in M_2$ but $w \notin M_2$, we must have $wv_{22} \in E(G)$ because M_2 is a module. Then since $w, u \in M_1$ and $v_{22} \notin M_1$ we must have $uv_{22} \in E(G)$ because M_1 is a module. A symmetric argument proves that $uv_{22} \in E(G)$ implies $wu \in E(G)$. It follows that $\mathbf{A}_G(u, w) = \mathbf{A}_G(u, v_{22})$. So again $\mathbf{A}_G[U, W \cup M_{11} \cup M_{22}]$ can be obtained from $\mathbf{A}_G[U, W]$ by copying columns, and thus the two matrices have the same rank.

Since β is bijective, this proves that the rank of any bipartite adjacency matrix induced by removing an edge $e \in T$ is bounded by d . We conclude that the width of \mathcal{T} is at most d and thus $rw(G[M_1 \cup M_2]) \leq d$. \square

Definition 16. Let G be a graph and $d \in \mathbb{N}$. We define a relation \sim_d^G on $V(G)$ by letting $v \sim_d^G w$ if and only if there is a module M of G with $v, w \in M$ and $rw(G[M]) \leq d$. We drop the superscript from \sim_d^G if the graph G is clear from context.

Proposition 17. For every graph G and $d \in \mathbb{N}$ the relation \sim_d is an equivalence relation, and each equivalence class U of \sim_d is a module of G with $rw(G[U]) \leq d$.

Proof. Let G be a graph and $d \in \mathbb{N}$. For every $v \in V(G)$, the singleton $\{v\}$ is a module of G , so \sim_d is reflexive. Symmetry of \sim_d is trivial. For transitivity, let $u, v, w \in V(G)$ such that $u \sim_d v$ and $v \sim_d w$. Then there are modules M_1, M_2 of G such that $u, v \in M_1$, $v, w \in M_2$, and $rw(G[M_1]), rw(G[M_2]) \leq d$. By Lemma 15 $M_1 \cup M_2$ is a module of G with $rw(G[M_1 \cup M_2]) \leq d$. In combination with $u, w \in M_1 \cup M_2$ that implies $u \sim_d w$. This concludes the proof that \sim_d is an equivalence relation.

Now let $v \in V(G)$ and let $U = [v]_{\sim_d}$. For each $u \in U$ there is a module W_u of G with $u, v \in W_u$ and $rw(G[W_u]) \leq d$. By Lemma 15, $W = \bigcup_{u \in U} W_u$ is a module of G and $rw(G[W]) \leq d$. Clearly, $[v]_{\sim_d} \subseteq W$. On the other hand, $u \in W$ implies $v \sim_d u$ by definition of \sim_d , so $W \subseteq [v]_{\sim_d}$. That is, $W = [v]_{\sim_d}$. \square

Corollary 18. Let G be a graph and $d \in \mathbb{N}$. The equivalence classes of \sim_d form a smallest rank-width- d cover of G .

Proof. Let $\mathcal{U} = \{U_1, \dots, U_k\}$ be the set of equivalence classes of \sim_d . It is immediate from Proposition 17 that \mathcal{U} is a rank-width- d cover of G . Let $\mathcal{V} = \{V_1, \dots, V_j\}$ be a partition of $V(G)$ with $j < k$. By the pigeonhole principle, there have to be vertices $v_1, v_2 \in V(G)$ and indices $i_1, i_2 \in [k]$, $i \in [j]$ such that $v_1, v_2 \in V_j$ but $v_1 \in U_{i_1}$ and $v_2 \in U_{i_2}$, where $i_1 \neq i_2$. Thus $v_1 \not\sim_d v_2$, so there is no module M of $V(G)$ such that $v_1, v_2 \in M$ and $rw(G[M]) \leq d$. In particular, V_i is not a module or $rw(G[V_i]) > d$. So \mathcal{V} is not a rank-width- d cover of G . \square

Proposition 19. Let $d \in \mathbb{N}$ be a constant. Given a graph G and two vertices $v, w \in V(G)$, we can decide whether $v \sim_d w$ in polynomial time.

Proof. By Lemma 14 we can compute the unique minimal (with respect to set inclusion) module M containing v and w in time $O(|V(G)|^2)$. Since rank-width is preserved for induced subgraphs, there is a module M' containing v and w with $rw(G[M']) \leq d$ if and only if $rw(G[M]) \leq d$. By Theorem 9 this can be decided in time $O(|V(G)|^3)$. \square

Proof (of Theorem 1). Let $d \in \mathbb{N}$ be a constant. Given a graph G , we can compute the set of equivalence classes of \sim_d by testing whether $v \sim_d w$ for each pair of vertices $v, w \in V(G)$. By Proposition 19, this can be done in polynomial time, and by Corollary 18, $V(G)/\sim_d$ is a smallest rank-width- d cover of G . \square

4 Kernels for MSO Model Checking

In this section, we show that every MSO model checking problem admits a polynomial kernel when parameterized by the \mathcal{C} -cover number of the input graph, where \mathcal{C} is some recursively enumerable class of graphs satisfying the following properties:

- (I) \mathcal{C} contains the single-vertex graph, and a \mathcal{C} -cover of a graph G with minimum cardinality can be computed in polynomial time.
- (II) There is an algorithm \mathbb{A} that decides whether $G \models \varphi$ in time polynomial in $|V(G)|$ for any fixed MSO sentence φ and any graph $G \in \mathcal{C}$.

Let G be a graph and $U \subseteq V(G)$. Let \mathbf{v} be an m -tuple of vertices of G , and let \mathbf{V} be an l -tuple of sets of vertices of G . We write $\mathbf{V}|_U = (V_1 \cap U, \dots, V_l \cap U)$ to refer to the elementwise intersection of \mathbf{V} with U . Similarly, we let $\mathbf{v}|_U = (v_{i_1}, \dots, v_{i_t})$, $t \leq m$ denote the subsequence of elements from \mathbf{v} contained in U . If $\{U_1, \dots, U_k\}$ is a modular partition of G and $i \in [k]$ we will abuse notation and write $\mathbf{v}|_i = \mathbf{v}|_{U_i}$ and $\mathbf{V}|_i = \mathbf{V}|_{U_i}$ if there is no ambiguity about what partition the index belongs to.

Definition 20 (Congruent). Let q and l be non-negative integers and let G and G' be graphs with modular partitions $\{M_1, \dots, M_k\}$ and $\{M'_1, \dots, M'_k\}$, respectively. Let \mathbf{V}_0 be an l -tuple of subsets of $V(G)$ and let \mathbf{U}_0 be an l -tuple of subsets of $V(G')$. We say $(G, \mathbf{M}, \mathbf{V}_0)$ and $(G', \mathbf{M}', \mathbf{U}_0)$ are q -congruent if the following conditions are met:

1. For every $i, j \in [k]$ with $i \neq j$, M_i and M_j are adjacent in G if and only if M'_i and M'_j are adjacent in G' .
2. For each $i \in [k]$, $\text{type}_q(G[M_i], \mathbf{V}_0|_i) = \text{type}_q(G'[M'_i], \mathbf{U}_0|_i)$

Lemma 21. Let q and l be non-negative integers and let G and G' be graphs with modular partitions $\{M_1, \dots, M_k\}$ and $\{M'_1, \dots, M'_k\}$. Let \mathbf{V}_0 be an l -tuple of subsets of $V(G)$ and let \mathbf{U}_0 be an l -tuple of subsets of $V(G')$. If $(G, \mathbf{M}, \mathbf{V}_0)$ and $(G', \mathbf{M}', \mathbf{U}_0)$ are q -congruent, then $\text{type}_q(G, \mathbf{V}_0) = \text{type}_q(G', \mathbf{U}_0)$.

Proof. For $i \in [k]$, we write $G_i = G[M_i]$ and $G'_i = G'[M'_i]$. By Theorem 8, Condition 2 of Definition 20 is equivalent to $(G_i, \mathbf{V}_0|_i) \equiv_q^{\text{MSO}} (G'_i, \mathbf{U}_0|_i)$. That is, for each $i \in [k]$, duplicator has a winning strategy π_i in the q -round MSO game played on G_i and G'_i starting from $(\mathbf{V}_0|_i, \mathbf{U}_0|_i)$. We construct a strategy witnessing $(G, \mathbf{V}_0) \equiv_q^{\text{MSO}} (G', \mathbf{U}_0)$ by aggregating duplicator's moves from these k games in the following way:

1. Suppose spoiler makes a set move W and assume without loss of generality that $W \subseteq V(G)$. For $i \in [k]$, let $W_i = M_i \cap W$, and let W'_i be duplicator's response to W_i according to π_i . Then duplicator responds with $W' = \cup_{i=1}^k W'_i$.
2. Suppose spoiler makes a point move s and again assume without loss of generality that $s \in V(G)$. Then $s \in M_i$ for some $i \in [k]$. Duplicator responds with $s' \in M'_i$ according to π_i .

Assume duplicator plays according to this strategy and consider a play of the q -round MSO game on G and G' starting from $(\mathbf{V}_0, \mathbf{U}_0)$. Let $v_1, \dots, v_m \in V(G)$ and $u_1, \dots, u_m \in V(G')$ be the point moves and $V_1, \dots, V_l \subseteq V(G)$ and $U_1, \dots, U_l \subseteq V(G')$ be the set moves, so that $l + m = q$ and the moves made in the same round have the same index. We claim that (\mathbf{v}, \mathbf{u}) defines a partial isomorphism between $(G, \mathbf{V}_0 \mathbf{V})$ and $(G', \mathbf{U}_0 \mathbf{U})$.

- Let $j_1, j_2 \in [m]$ and let $i_1, i_2 \in [k]$ such that $v_{j_1} \in M_{i_1}$ and $v_{j_2} \in M_{i_2}$. Suppose $i_1 = i_2 = i$. Since duplicator plays according to a winning strategy in the game on G_i and G'_i , the restriction $(\mathbf{v}|_i, \mathbf{u}|_i)$ defines a partial isomorphism between $(G_i, (\mathbf{V}_0 \mathbf{V})|_i)$ and $(G'_i, (\mathbf{U}_0 \mathbf{U})|_i)$. It follows that $(v_{j_1}, v_{j_2}) \in E(G)$ if and only if $(u_{j_1}, u_{j_2}) \in E(G')$ and $v_{j_1} = v_{j_2}$ if and only if $u_{j_1} = u_{j_2}$. Now suppose $i_1 \neq i_2$. Then $v_{j_1} \neq v_{j_2}$ and also $u_{j_1} \neq u_{j_2}$ since $u_{j_1} \in M'_{i_1}$ and $u_{j_2} \in M'_{i_2}$ by choice of duplicator's strategy. By congruence, M_{i_1} and M_{i_2} are adjacent in G if and only if M'_{i_1} and M'_{i_2} are adjacent in G' , so we must have $(v_{j_1}, v_{j_2}) \in E(G)$ if and only if $(u_{j_1}, u_{j_2}) \in E(G')$.
- Let $j \in [m]$ and let $i \in [k]$ such that $v_j \in M_i$. By construction of duplicator's strategy, we have $u_j \in M'_i$. Note that if $x \in S$ then $x \in S'$ if and only if $x \in S'|_S$ for arbitrary sets S and S' . Combined with the fact that $(\mathbf{v}|_i, \mathbf{u}|_i)$ defines a partial isomorphism between $(G_i, (\mathbf{V}_0 \mathbf{V})|_i)$ and $(G'_i, (\mathbf{U}_0 \mathbf{U})|_i)$, this observation implies that v_i is contained in any of the sets from $\mathbf{V}_0 \mathbf{V}$ if and only if u_i is contained in the sets from $\mathbf{U}_0 \mathbf{U}$ with the same indices.

□

Lemma 22. Let \mathcal{C} be a recursively enumerable graph class and let q be a non-negative integer constant. Let $G \in \mathcal{C}$ be a graph. If $G \models \varphi$ can be decided in time polynomial in $|V(G)|$ for any fixed $\varphi \in \text{MSO}_{q,0}$ then one can in polynomial time compute a graph $G' \in \mathcal{C}$ such that $|V(G')|$ is bounded by a constant and $\text{type}_q(G) = \text{type}_q(G')$.

Proof. By Lemma 5 we can compute a formula Φ capturing the type T of G in polynomial time. Given Φ , a graph $G' \in \mathcal{C}$ satisfying Φ can be effectively computed as follows. We start enumerating \mathcal{C} and check for each graph $G' \in \mathcal{C}$ whether $G' \models \Phi$. If this is the case, we stop and output G' . Since $G \models \Phi$ this procedure must terminate eventually. Fixing \mathcal{C} and the order in which graphs are enumerated, the number of graphs we have to check depends only on T . By Fact 1 the number of rank q -types is finite for each q , so we can think of the total number of checks as bounded by a constant. Moreover the time spent on each check depends only on T and the size of the graph G' . Because the number of graphs enumerated is bounded by a constant, we can think of the latter as bounded by a constant as well. Thus the algorithm computing a model of Φ runs in constant time. \square

Lemma 23. *Let q be a non-negative integer constant, and let \mathcal{C} be a recursively enumerable graph class satisfying (II). Then given a graph G and a \mathcal{C} -cover $\{U_1, \dots, U_k\}$, one can in polynomial time compute a graph G' with modular partition $\{U'_1, \dots, U'_k\}$ such that (G, \mathbf{U}) and (G', \mathbf{U}') are q -congruent and for each $i \in [k]$, $G'[U'_i] \in \mathcal{C}$ and the number of vertices in U'_i is bounded by a constant.*

Proof. For each $i \in [k]$, we compute a graph $G'_i \in \mathcal{C}$ of constant size with the same MSO rank- q type as $G_i = G[U_i]$. By Lemma 22, this can be done in polynomial time. Now let G' be the graph obtained from the disjoint union of the graphs G'_i for $i \in [k]$ as follows. For $i \in [k]$, let U'_i denote the set of vertices from the copy of G'_i . If U_i and U_j are adjacent in G for $i, j \in [k]$ and $i \neq j$, we insert an edge vw for every $v \in U'_i$ and $w \in U'_j$. Then U'_1, \dots, U'_k is a modular partition of G' , and for $i, j \in [k]$ and $i \neq j$, modules U_i and U_j are adjacent in G if and only if U'_i and U'_j are adjacent in G' . It is readily verified that (G, \mathbf{U}) and (G', \mathbf{U}') are q -congruent. \square

Proposition 24. *Let φ be a fixed MSO sentence. Let \mathcal{C} be a recursively enumerable graph class satisfying (I) and (II). Then MSO-MC_φ has a polynomial kernel parameterized by the \mathcal{C} -cover number of the input graph.*

Proof. Let G be a graph with \mathcal{C} -cover number k , and let $\{U_1, \dots, U_k\}$ be a smallest \mathcal{C} -cover given by (I). Let q be the quantifier rank of φ . By Lemma 23 and (II), we can in polynomial time compute a graph G' and a modular partition $\{U'_1, \dots, U'_k\}$ of G' such that (G, \mathbf{U}) and (G', \mathbf{U}') are q -congruent and for each $i \in [k]$, $|U'_i|$ is bounded by a constant. It follows from Lemma 21 that $\text{type}_q(G) = \text{type}_q(G')$. In particular, $G \models \varphi$ if and only if $G' \models \varphi$. Moreover, we have $|V(G')| \in O(k)$, so G' is a polynomial kernel. \square

Proof (of Theorem 2). Immediate from Theorems 1, 9, and 10 in combination with Proposition 24. \square

Corollary 25. *The following problems have polynomial kernels when parameterized by the rank-width- d cover number of the input graph: INDEPENDENT DOMINATING SET, c -COLORING, c -DOMATIC NUMBER, c -PARTITION INTO TREES, c -CLIQUE COVER, c -PARTITION INTO PERFECT MATCHINGS, c -COVERING BY COMPLETE BIPARTITE SUBGRAPHS.*

5 Kernels for MSO Optimization

By definition, MSO formulas can only directly capture decision problems such as 3-colorability, but many problems of interest are formulated as optimization problems. The usual way of transforming decision problems into optimization problems does not work here, since the MSO language cannot handle arbitrary numbers.

Nevertheless, there is a known solution. Arnborg, Lagergren, and Seese [2] (while studying graphs of bounded tree-width), and later Courcelle, Makowsky, and Rotics [9] (for graphs of bounded clique-width), specifically extended the expressive power of MSO logic to define so-called LINEMS optimization problems, and consequently showed the existence of efficient (parameterized) algorithms for such problems in the respective cases.

The MSO optimization problems (problems of the form $\text{MSO-OPT}_\varphi^\diamond$) considered here are a streamlined and simplified version of the formalism introduced in [9]. Specifically, we consider only a single free variable X , and ask for a satisfying assignment of X with minimum or maximum cardinality. To achieve our results, we need a recursively enumerable graph class \mathcal{C} that satisfies (I) and (II) along with the following property:

- (III) Let $\varphi = \varphi(X)$ be a fixed MSO formula. Given a graph $G \in \mathcal{C}$, a set $S \subseteq V(G)$ of minimum (maximum) cardinality such that $G \models \varphi(S)$ can be found in polynomial time, if one exists.

Our approach will be similar to the MSO kernelization algorithm, with one key difference: when replacing the subgraph induced by a module, the cardinalities of subsets of a given q -type may change, so we need to keep track of their cardinalities in the original subgraph.

To do this, we introduce an annotated version of $\text{MSO-}\text{OPT}_\varphi^\diamond$. Given a graph $G = (V, E)$, an *annotation* \mathcal{W} is a set of triples (X, Y, w) with $X \subseteq V, Y \subseteq V, w \in \mathbb{N}$. For every set $Z \subseteq V$ we define

$$\mathcal{W}(Z) = \sum_{(X, Y, w) \in \mathcal{W}, X \subseteq Z, Y \cap Z = \emptyset} w.$$

We call the pair (G, \mathcal{W}) an *annotated graph*. If the integer w is represented in binary, we can represent a triple (X, Y, w) in space $|X| + |Y| + \log_2(w)$. Consequently, we may assume that the size of the encoding of an annotated graph (G, \mathcal{W}) is polynomial in $|V(G)| + |\mathcal{W}| + \max_{(X, Y, w) \in \mathcal{W}} \log_2 w$.

Each MSO formula $\varphi(X)$ and $\diamond \in \{\leq, \geq\}$ gives rise to an *annotated MSO-optimization problem*.

$a\text{MSO-}\text{OPT}_\varphi^\diamond$

Instance: A graph G with an annotation \mathcal{W} and an integer $r \in \mathbb{N}$.

Question: Is there a set $Z \subseteq V(G)$ such that $G \models \varphi(Z)$ and $\mathcal{W}(Z) \diamond r$?

Notice that any instance of $\text{MSO-}\text{OPT}_\varphi^\diamond$ is also an instance of $a\text{MSO-}\text{OPT}_\varphi^\diamond$ with the trivial annotation $\mathcal{W} = \{(\{v\}, \emptyset, 1) : v \in V(G)\}$. The main result of this section is a bikernelization algorithm which transforms any instance of $\text{MSO-}\text{OPT}_\varphi^\diamond$ into an instance of $a\text{MSO-}\text{OPT}_\varphi^\diamond$; this kind of bikernel is called an *annotated kernel* [1].

The results below are stated and proved for minimization problems $a\text{MSO-}\text{OPT}_\varphi^\leq$ only. This is without loss of generality – the proofs for maximization problems are symmetric.

Lemma 26. *Let q and l be non-negative integers and let G and G' be graphs such that G and G' have the same $q + l$ MSO type. Then for any l -tuple \mathbf{V} of sets of vertices of G , there exists an l -tuple \mathbf{U} of sets of vertices of G' such that $\text{type}_q(G, \mathbf{V}) = \text{type}_q(G', \mathbf{U})$.*

Proof. Suppose there exists an l -tuple \mathbf{V} of sets of vertices of G , and a formula $\varphi = \varphi(X_1, \dots, X_l) \in \text{MSO}_{q,l}$ such that $G \models \varphi(V_1, \dots, V_l)$ but for every l -tuple \mathbf{U} of sets of vertices of G' we have $G' \not\models \varphi(U_1, \dots, U_l)$. Let $\psi = \exists X_1 \dots \exists X_l \varphi$. Clearly, $\psi \in \text{MSO}_{q+l,0}$ and $G \models \psi$ but $G' \not\models \psi$, a contradiction. \square

Lemma 27. *Let $\varphi = \varphi(X)$ be a fixed MSO formula and \mathcal{C} be a recursively enumerable graph class satisfying (II) and (III). Then given an instance (G, r) of $\text{MSO-}\text{OPT}_\varphi^\leq$ and a \mathcal{C} -cover $\{U_1, \dots, U_k\}$ of G , an annotated graph (G', \mathcal{W}) satisfying the following properties can be computed in polynomial time.*

1. $(G, r) \in \text{MSO-}\text{OPT}_\varphi^\leq$ if and only if $(G', \mathcal{W}, r) \in a\text{MSO-}\text{OPT}_\varphi^\leq$.
2. $|V(G')| \in O(k)$.
3. The encoding size of (G', \mathcal{W}) is $O(k \log(|V(G)|))$.

Proof. Let q be the quantifier rank of φ . By Lemma 23 and (II), we can in polynomial time compute a graph G' and a modular partition $\{U'_1, \dots, U'_k\}$ of G' such that (G, \mathbf{U}) and (G', \mathbf{U}') are $(q + 1)$ -congruent, $|U'_i|$ is bounded by a constant, and $G'[U'_i] \in \mathcal{C}$ for each $i \in [k]$. To compute the annotation \mathcal{W} , we proceed as follows. For each $i \in [k]$, we go through all subsets $W' \subseteq U'_i$. By Lemma 5, we can compute a formula Φ such that for any graph H and $W \subseteq V(H)$ we have $\text{type}_q(G'[U'_i], W) = \text{type}_q(H, W)$ if and only if $H \models \Phi(W)$. Since $|U'_i|$ has constant size for every $i \in [k]$, this can be done within a constant time bound. By Lemma 26 and because (G, \mathbf{U}) and (G', \mathbf{U}') are $(q + 1)$ -congruent, there has to be a $W \subseteq U_i$ such that $G_i \models \Phi(W)$. Using the algorithm given by (III), we can compute a minimum-cardinality subset $W^* \subseteq U_i$ with this property in polynomial time. We then add the triple $(W', U'_i \setminus W', |W^*|)$ to \mathcal{W} . In total, the number of subsets processed is in $O(k)$. From this observation we get the desired bounds on the total runtime, $|V(G')|$, and the encoding size of (G', \mathcal{W}) .

We claim that $(G', \mathcal{W}, r) \in a\text{MSO-}\text{OPT}_\varphi^\leq$ if and only if $(G, r) \in \text{MSO-}\text{OPT}_\varphi^\leq$. Suppose there is a set $W \subseteq V(G)$ of vertices such that $G \models \varphi(W)$ and $|W| \leq r$. Since U_1, \dots, U_k is a partition of $V(G)$, we have $W = \cup_{i \in [k]} W_i$, where $W_i = W \cap U_i$. For each $i \in [k]$, let $W_i^* \subseteq U_i$ be a subset of minimum cardinality such that $\text{type}_q(G[U_i], W_i) = \text{type}_q(G[U_i], W_i^*)$. By Lemma 26 and $(q + 1)$ -congruence of (G, \mathbf{U}) and (G', \mathbf{U}') , there

is $W'_i \subseteq U'_i$ for each $i \in [k]$ such that $\text{type}_q(G'[U'_i], W'_i) = \text{type}_q(G[U_i], W_i^*)$. By construction, \mathcal{W} contains a triple $(W'_i, U'_i \setminus W'_i, |W_i^*|)$. Observe that $(X, Y, w) \in \mathcal{W}$ and $(X, Y, w') \in \mathcal{W}$ implies $w = w'$. Let $W' = \bigcup_{i \in [k]} W'_i$. Then by $(q+1)$ -congruence of (G, U) and (G', U') and Lemma 21, we must have $\text{type}_q(G, W) = \text{type}_q(G', W')$. In particular, $G' \models \varphi(W')$. Furthermore,

$$\mathcal{W}(W') = \sum_{(W'_i, U'_i \setminus W'_i, |W_i^*|) \in \mathcal{W}, U'_i \cap W' = W'_i} |W_i^*| \leq \sum_{i \in [k]} |W_i| = |W| \leq r.$$

For the converse, let $W' \subseteq V(G')$ such that $\mathcal{W}(W') \leq r$ and $G' \models \varphi(W')$, let W'_i denote $W' \cap U'_i$ for $i \in [k]$. By construction, there is a set $W_i \subseteq U_i$ for each $i \in [k]$ such that $\text{type}_q(G[U_i], W_i) = \text{type}_q(G'[U'_i], W'_i)$ and $\mathcal{W}(W') = \sum_{i \in [k]} |W_i|$. Let $W = \bigcup_{i \in [k]} W_i$. Then by congruence and Lemma 21 we get $\text{type}_q(G, W) = \text{type}_q(G', W')$ and thus $G \models \varphi(W)$. Moreover, $|W| = \mathcal{W}(W') \leq r$. \square

Fact 2 (Folklore). *Given an MSO sentence φ and a graph G , one can decide whether $G \models \varphi$ in time $O(2^{nl})$, where $n = |V(G)|$ and $l = |\varphi|$.*

Proposition 28. *Let $\varphi = \varphi(X)$ be a fixed MSO formula, and let \mathcal{C} be a recursively enumerable graph class satisfying (I), (II), and (III). Then $\text{MSO-OPT}_{\varphi}^{\leq}$ has a polynomial bikernel parameterized by the \mathcal{C} -cover number of the input graph.*

Proof. Let (G, r) be an instance of $\text{MSO-OPT}_{\varphi}^{\leq}$. By (I) a smallest \mathcal{C} -cover $\{U_1, \dots, U_k\}$ of G can be computed in polynomial time. Let (G', \mathcal{W}) be an annotated graph computed from G and $\{U_1, \dots, U_k\}$ according to Lemma 27. Let $n = |V(G)|$ and suppose $2^k \leq n$. Then we can solve (G', \mathcal{W}, r) in time n^c for some constant c that only depends on φ and \mathcal{C} . To do this, we go through all $2^{O(k)}$ subsets W of G' and test whether $\mathcal{W}(W) \leq r$. If that is the case, we check whether $G' \models \varphi(W)$. By Fact 2 this check can be carried out in time $c_1 2^{c_2 k} \leq c_1 n^{c_2}$ for suitable constants c_1 and c_2 depending only on \mathcal{C} and φ . Thus we can find a c such that the entire procedure runs in time n^c whenever n is large enough. If we find a solution $W \subseteq V(G')$ we return a trivial yes-instance; otherwise, a trivial no-instance (of $\text{aMSO-OPT}_{\varphi}^{\leq}$). Now suppose $n < 2^k$. Then $\log(n) < k$ and so the encoding size of \mathcal{W} is polynomial in k . Thus (G', \mathcal{W}, r) is a polynomial bikernel. \square

Proof (of Theorem 3). Immediate from Theorems 1, 9, and 10 when combined with Proposition 28. \square

Corollary 29. *The following problems have polynomial bikernels when parameterized by the rank-width- d cover number of the input graph: MINIMUM DOMINATING SET, MINIMUM VERTEX COVER, MINIMUM FEEDBACK VERTEX SET, MAXIMUM INDEPENDENT SET, MAXIMUM CLIQUE, LONGEST INDUCED PATH, MAXIMUM BIPARTITE SUBGRAPH, MINIMUM CONNECTED DOMINATING SET.*

6 Conclusion

Recently Bodlaender et al. [4] and Fomin et al. [16] established *meta-kernelization theorems* that provide polynomial kernels for large classes of parameterized problems. The known meta-kernelization theorems apply to optimization problems parameterized by *solution size*. Our results are, along with very recent results parameterized by the modulator to constant-treedepth [17], the first meta-kernelization theorems that use a *structural parameter* of the input and not the solution size. In particular, we would like to emphasize that our Theorem 3 applies to a large class of optimization problems where the solution size can be arbitrarily large.

It is also worth noting that our structural parameter, the rank-width- d cover number, provides a trade-off between the maximum rank-width of modules (the constant d) and the maximum number of modules (the parameter k). Different problem inputs might be better suited for smaller d and larger k , others for larger d and smaller k . This two-dimensional setting could be seen as a contribution to a *multivariate complexity analysis* as advocated by Fellows et al. [13].

We conclude by mentioning possible directions for future research. We believe that some of our results can be extended from modular partitions to partitions into splits [7].¹ This would indeed result in a more general parameter, however the precise details would still require further work (one problem is that while all modules are partitive, only strong splits have this property). Another direction would then be to focus on polynomial kernels for problems which cannot be described by MSO logic, such as HAMILTONIAN PATH or CHROMATIC NUMBER.

¹ We thank Sang-il Oum for pointing this out to us.

References

1. F. N. Abu-Khzam and H. Fernau. Kernels: Annotated, proper and induced. In *Parameterized and Exact Computation*, Lecture Notes in Computer Science, pages 264–275. Springer Verlag, 2006.
2. S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *J. Algorithms*, 12(2):308–340, 1991.
3. H. L. Bodlaender, R. G. Downey, M. R. Fellows, and D. Hermelin. On problems without polynomial kernels. *J. of Computer and System Sciences*, 75(8):423–434, 2009.
4. H. L. Bodlaender, F. V. Fomin, D. Lokshtanov, E. Penninkx, S. Saurabh, and D. M. Thilikos. (meta) kernelization. In *FOCS 2009*, pages 629–638. IEEE Computer Society, 2009.
5. B.-M. Bui-Xuan, M. Habib, V. Limouzy, and F. de Montgolfier. Algorithmic aspects of a general modular decomposition theory. *Discr. Appl. Math.*, 157(9):1993–2009, 2009.
6. B.-M. Bui-Xuan, J. A. Telle, and M. Vatshelle. Boolean-width of graphs. *Theoretical Computer Science*, 412(39):5187–5204, 2011.
7. P. Charbit, F. de Montgolfier, and M. Raffinot. Linear time split decomposition revisited. *SIAM J. Discrete Math.*, 26(2):499–514, 2012.
8. M. Chein, M. Habib, and M. Maurer. Partitive hypergraphs. *Discrete Math.*, 37(1):35–50, 1981.
9. B. Courcelle, J. A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000.
10. R. Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer Verlag, New York, 2nd edition, 2000.
11. R. Downey, M. R. Fellows, and U. Stege. Parameterized complexity: A framework for systematically confronting computational intractability. In *Contemporary Trends in Discrete Mathematics: From DIMACS and DIMATIA to the Future*, volume 49 of *AMS-DIMACS*, pages 49–99. American Mathematical Society, 1999.
12. M. R. Fellows. The lost continent of polynomial time: Preprocessing and kernelization. In *IWPEC 2006*, volume 4169 of *Lecture Notes in Computer Science*, pages 276–277. Springer Verlag, 2006.
13. M. R. Fellows, B. M. Jansen, and F. Rosamond. Towards fully multivariate algorithmics: Parameter ecology and the deconstruction of computational complexity. *European J. Combin.*, 34(3):541–566, 2013.
14. J. Flum and M. Grohe. *Parameterized Complexity Theory*, volume XIV of *Texts in Theoretical Computer Science. An EATCS Series*. Springer Verlag, Berlin, 2006.
15. F. V. Fomin. Kernelization. In *CSR 2010*, volume 6072 of *Lecture Notes in Computer Science*, pages 107–108. Springer Verlag, 2010.
16. F. V. Fomin, D. Lokshtanov, N. Misra, and S. Saurabh. Planar f-deletion: Approximation, kernelization and optimal fpt algorithms. In *FOCS 2012*, pages 470–479. IEEE Computer Society, 2012.
17. J. Gajarský, P. Hliněný, J. Obdržálek, S. Ordyniak, F. Reidl, P. Rossmanith, F. S. Villaamil, and S. Sikdar. Kernelization using structural parameters on sparse graph classes. *CoRR*, abs/1302.6863, 2013.
18. R. Ganian and P. Hliněný. On parse trees and Myhill-Nerode-type tools for handling graphs of bounded rank-width. *Discr. Appl. Math.*, 158(7):851–867, 2010.
19. J. Guo and R. Niedermeier. Invitation to data reduction and problem kernelization. *ACM SIGACT News*, 38(2):31–45, Mar. 2007.
20. P. Hliněný and S. Oum. Finding branch-decompositions and rank-decompositions. *SIAM J. Comput.*, 38(3):1012–1032, 2008.
21. M. Lampis. Algorithmic meta-theorems for restrictions of treewidth. *Algorithmica*, 64(1):190–37, 2012.
22. L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
23. N. Misra, V. Raman, and S. Saurabh. Lower bounds on kernelization. *Discrete Optimization*, 8(1):110–128, 2011.
24. S. Oum and P. Seymour. Approximating clique-width and branch-width. *J. Combin. Theory Ser. B*, 96(4):514–528, 2006.
25. F. Rosamond. Table of races. In *Parameterized Complexity Newsletter*, pages 4–5. 2010. <http://fpt.wikidot.com/>.